



Distributed Systems Part II

Exercise Sheet 5

1 Paxos Timeline

In the following, we want to run an implementation of Paxos on a set of 3 nodes (A, B, C) that act as acceptors. Moreover, we assume that there are two more nodes (Q, R) that act as proposers. The implementation of the acceptors is exactly as shown in the lecture, Slide 49. The two proposers use the implementation given in Figure 1.

```
function suggestValue(Node N1, Node N2, Timeout t, Value x, RequestNumber n)

  Send prepare(x,n) to nodes N1, N2
  Wait t seconds
  if within these t seconds, either N1 or N2 has not replied then
    suggestValue(N1, N2, t, x, n + 2)

  Let (y, m) be the received proposal with the largest request number
  if m == 0 then
    u := x
  else
    u := y
  Send propose(u, n) to N1, N2
  Wait t seconds
  if within these t seconds, either N1 or N2 has not replied with ack(u, n) then
    suggestValue(N1, N2, t, x, n + 2)

  Print("value is chosen: " + u)
```

Figure 1: Code executed by the proposers.

Draw a timeline containing all transmitted messages if a user invokes `suggestValue(A, B, 1, 22, 1)` on Q at time T_0 and `suggestValue(B, C, 2, 33, 2)` on R at time $T_0 + 0.5sec$! We assume that processing times on the nodes can be neglected (i.e. is zero), and that all messages arrive within less than $0.5sec$.

2 Paxos Acceptors

In the lecture you have seen how Paxos can solve consensus without the need of a single coordinator. It lets each node execute one or more of the following roles: proposer, acceptor, and learner. In this task you will have a closer look at the purpose of the acceptor.

- Assume, that in a network of 5 nodes there is one node with a faulty register that is used to store the value n_{max} of its acceptor programm (See slide *Paxos: Algorithm of the Acceptor*). Can this pose a problem to the Paxos algorithm? Explain what happens in the worst case scenario.
- Paxos assumes that there are no Byzantine failures. That is, memory failures such as the one described above, are not assumed to happen. Under these original assumptions, can you explain what the purpose of the *prepare step* in Paxos is?